DOCUMENT RESUME

ED 161 086 CS 204 489

AUTHOR Miller, James R.; Hayes-Roth, Barbara
TITLE Text Annotation: A Technique for Facilitating Knowledge Integration.
INSTITUTION Rand Corp., Santa Monica, Calif.
PUB DATE Dec 77
NOTE 22p.
AVAILABLE FROM The Rand Corporation, Santa Monica, California 90406 ($1.50)

EDRS PRICE MF-$0.83 HC-$1.67 Plus Postage.
DESCRIPTORS *Cognitive Processes; Comprehension; Higher Education; *Learning Processes; *Memory; Psychological Studies; *Reading Research; Retention; Thought Processes
IDENTIFIERS *Annotation; *Inference; Prose Learning

ABSTRACT

Thirty-six undergraduate students participated in a two-hour experiment that measured the effect of fact and inference text annotation on the integration of information. The subjects were asked to read four pairs of meaningful stories, each of which contained four pairs of related facts about a mythical country. The texts of the experimental group included two methods of facilitating inference verification: fact annotation which repeated the premises that occurred in the first text as footnotes to the appropriate premises in the second text, and fact-and-inference annotation which specified the appropriate inference as well as repeated the fact annotation. After a self-paced reading of the materials, subjects were tested for the immediate recognition of the sentences they had just read; twenty minutes later they were tested for inference and syllogism judgments. While both annotation methods improved the inference verification of the experimental group over that of the control group, fact annotation resulted in better overall performance. Fact-and-inference annotation produced incorrect inferences and limited the generation of inferences to those specified by the annotations. (MAI)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED161086

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

TEXT ANNOTATION: A TECHNIQUE FOR FACILITATING KNOWLEDGE INTEGRATION

James R. Miller and Barbara Hayes-Roth

December 1977

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

The Rand Corporation

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC) AND USERS OF THE ERIC SYSTEM."

CS204489

P-6054

The Rand Paper Series

Papers are issued by The Rand Corporation as a service to its professional Staff. Their purpose is to facilitate the exchange of ideas among those who share the author's research interests; Papers arr not reports prepared in fulfillment of Rand's contracts or grants. Views expressed in a Paper are the author's own, and are not necessarily shared by Rand or its research sponsors.

The Rand Corporation
Santa Monica, California 90406

ABSTRACT

Subjects read pairs of related texts and then verified true and false inferences derived from the information in the texts. Verification of a true inference required integration of two facts (premises) that had occurred in separate texts. Two methods of annotating texts were evaluated as methods of facilitating inference verification. Fact annotation involved repeating the premises that had occurred in the first text as footnotes to the appropriate premises in the second text. Fact-and-inference annotation involved specifying the appropriate inference as well as repeating premises from the first text as footnotes to premises in the second text. While both annotation methods improved inference verification over a control group, fact annotation resulted in better performance overall. Fact-and-inference annotation led to more false alarms to incorrect inferences, and the generation of fewer inferences not directly specified by annotations.

Recent work on prose comprehension and memory has emphasized the importance of the integration of information in a text. Simply recording the individual facts present in a text does not constitute comprehension, does not produce good memory for the information in the text, and does not support higher-order cognitive processes such as inferencing. In order to achieve good comprehension of and memory for a text, the student must integrate the individual facts in the text with one another and with previously learned facts. That is, the student must identify referential and other relationships between acquired facts, and use them as the basis for organizing those facts in memory (cf. Bransford and Franks, 1971; Frase, 1969, 1973, 1975; Haviland and Clark, 1974; Hayes-Roth, 1977; see also Chapter III).

Text integration has typically been investigated by having subjects study a text containing one or more pairs of facts that can be combined to produce an inference. Later, subjects are tested on their knowledge of these inferences. For instance, Frase (1969) presented subjects with brief stories like the following:

> "The Fundalas are outcasts from other tribes in Central Ugala. It is the custom in this part of the country to get rid of certain types of people. The hill people of Central Ugala are farmers. The upper highlands provide excellent soil for cultivation. The farmers of this country are peace loving, which is reflected in their artwork. The outcasts of Central Ugala are all hill people. There are about fifteen different tribes in this area. (p. 2)"

Later, subjects were tested for their knowledge of the inference:

The Fundalas are hill people, (1)

which follows from:

The Fundalas are outcasts from other tribes... (2)

and

The outcasts of Central Ugala are all hill people. (3)

Working in a similar paradigm, Hayes-Roth and Thorndyke (Experiment 5, Chapter III) presented subjects pairs of related stories, such as those presented on pages 92-93. Later, subjects were tested for knowledge of the inference:

Albert Profiro hated King Egbert. (4)

which is based on the facts:

King Egbert was a dictator (from the first story). (5)

and

Albert Profiro hated all dictators (from the second story). (6)

Despite the importance of information integration, people are not very good at it. Frase (1969) found that free recall of his stories included relatively few inferences, even when subjects were told explicitly to include as many inferences as possible in their recall. A related result involves the number of inferences to which a subject is exposed. If an inference chain is represented by the linear order: A -> B -> C, a subject verifing the inference "A -> C" during acquisition would be exposed to only that one inference. Subjects verifing the more complex inference "A -> E", represented by the chain: A -> B -> C -> D -> E, should be exposed to the intermediate inferences of A -> C, B -> D, C -> E, A -> D, and B -> E, as well as the explicit inference A -> E. One might expect the latter condition to produce more inferences on a subsequent recall test. This did not occur.

Unpublished research from our laboratory indicates similar difficulties in information integration. Subjects (UCLA undergraduates enrolled in an introductory psychology course) read four pairs of related stories (the Hayes-Roth and Thorndyke materials illustrated above). Four pairs of related sentences occurred in each pair of stories, with one sentence from each sentence pair occurring in each story. Each pair of sentences could be used to derive an inference, as illustrated in (4) - (6) above. After studying all four pairs of stories, the subjects received a verification test on the sixteen true inferences and sixteen distractors. Eight of the sixteen subjects performed at or below chance on this verification test (as determined by a high threshold correction for guessing; Green and Swets, 1966). Providing three repetitions of each story did little to help performance: seven of the sixteen subjects were still no better than chance. Performance on these inference items was conditionalized on correctly answering a syllogism test corresponding to each inference item (see the method section below). The observed poor performance was then not due to faulty reasoning, but rather to the subjects' failure to integrate the related facts on which the inferences were based.

As these studies illustrate, an important educational goal must be to facilitate subjects' integration of related facts and their formation of inferences supported by acquired facts. The present research investigated the facilitative effects of two types of text annotation. Fact annotation involved placing a footnote in studied texts that simply repeated previously studied, related facts. Fact-and-Inference annotation involved footnotes that not only

repeated previously studied facts, but also presented the appropriate inference. In a Control condition, texts were not annotated. In the experiment to be presented here, subjects were presented the pairs of stories used by Hayes-Roth and Thorndyke, and were told to study the stories (and the footnotes in the experimental conditions) carefully and to generate as many inferences from the stories as they could. Subjects were later given a verification test on the true inferences and an equal number of distractors.

The experimental conditions differ in their theoretical costs and benefits. The Fact-and-Inference condition has the advantage that subjects are explicitly told the information to be deduced; there is no danger of their failing to generate a conclusion or generating a false conclusion. Given examples in the text, they also may gain experience with the logic behind inference construction, so that they might go on to generate more inferences by themselves. The danger of the Fact-and-Inference condition is that a presented inference may not be processed to as deep a level of memory (Craik and Lockhart, 1972) as in the Fact condition, where subjects must generate the inference themselves. This semantic activity, required in the Facts condition, is commonly found to increase long-term retention (Hyde & Jenkins, 1969). In the Fact-and-Inference condition, readers may skim over the inference without processing it to a sufficiently "deep" level to allow future retrieval.

It could be argued that, by this logic, Control subjects should do best of all, since they must locate both of the inference components by themselves, in addition to generating the inference. However, the performance of control subjects should suffer because of

uncertainty about where inferences can be generated. Although successful generation of an inference should produce excellent retention in the control condition, it is unlikely that very many inferences will be generated by these subjects. Thus, while Control subjects should perform very well on tests of inferences that they generated during study, they should perform worse than either Fact or Fact-and-Inference subjects overall.

Subjects in the Fact-and-Inference condition may perform somewhat better than subjects in the Fact condition on an immediate test, since they have just read the inferences they are being asked to verify. However, subjects in the Fact condition are given enough information in the footnotes to enable them to generate most of the inferences, and the semantic processing underlying inference generation should produce strong memory for these inferences. Subjects in the Fact condition may therefore perform even better than subjects in the Fact-and-Inference condition on a delayed test.

METHOD

Subjects. Thirty-six UCLA undergraduates participated in the two-hour experiment. Subjects were either paid $6.00 or given course credit for their participation.

Materials. Four pairs of meaningful stories were taken from Hayes-Roth and Thorndyke (Chapter IV). All of the stories were about the mythical country Morinthia. Each individual story was about a different topic, but the stories within a pair were about related topics. The four pairs of topics were: (a) The First Morinthian

Revolution and The Imprisonment of Albert Profiro; (b) Religious Customs and Beliefs in Morinthia and Religious Overtones during the Fever Epidemic in Morinthia; (c) The Marriage of Princess Isadora, Successor to the Throne of Morinthia and The Romance Between Princess Mathilde and Basil; and (d) The Provincial Lifestyle in Morinthia and The Morinthian Home of the Caledian Ambassador.

Each pair of stories included four pairs of related facts, such as (5) and (6) above. Each pair of related facts contained the information necessary to support a particular inference not explicitly stated in either story (such as (4) above). Related stories and, in particular, the related facts within the stories had identical wordings of common information.

Design. All subjects received all four pairs of stories in an order counterbalanced across subjects. To insure high retention of the stories, subjects received three copies of each story, alternating between Story 1 and Story 2 of that pair. Subjects were randomly assigned to one of the following experimental conditions:

Control condition: Control subjects were told (via printed instructions in their test booklets) that the experiment concerned their ability to read, understand, and remember prose, that they would read a number of stories about an imaginary country, and later be asked questions about the stories. The subjects were further instructed to try to combine information presented in the stories into inferences, and to write down any inferences they made in a space provided below the stories. An example (unrelated to any of the experimental stories) was presented, showing two stories with

information suitable for generating inferences. Subjects were told that generating inferences would maximize their performance on the tests to come, so they should try very hard to generate inferences.

Fact condition: Fact subjects received the same instructions and test materials as the control subjects, with the following additions. The second story of each pair was modified by placing a numerical footnote after each sentence in the story that was involved in one of the experimental inferences. This footnote referred the subject to a "hint" at the bottom of the page, which was actually a related sentence from the first story in the pair. Each of these pairs of sentences provided the basis for generating an inference. For example, referring back to the Hayes-Roth and Thorndyke texts above, the fact King Egbert was a dictator, which had occurred in the first story, would be repeated as a footnote to the fact Albert hated all dictators in the second story. The subjects were told that these hints could be used to make inferences, and that they should do so. They were also told that other inferences might be possible, so they should not restrict their inferences to those associated with hints. Again, subjects were told to write down all inferences they generated in a space provided on the page.

Fact-and-Inference condition: Fact-and-Inference subjects received the same materials as Fact subjects, with the following modifications. The second sentence of each pair was again marked with a footnote, but the hint indicated by the footnote was the related sentence from the first story in the pair and the inference that could be generated from these sentences. For example, both King Egbert was

a dictator and Therefore, Albert hated King Egbert would be given as a footnote to Albert hated all dictators. The subjects were told that these hints would show them examples of inferences between the two stories, and that they should learn these. They were also told to generate as many other inferences as possible, and write them down in the space provided.

Procedure. Subjects were tested in groups. Each subject was given a booklet containing the experimental stories and tests. Subjects' progress through the booklets was self-paced. Intentional learning instructions were given, including the warning that an inference test would be given.

Immediate recognition test: Immediately after all the stories of all four story pairs had been read, subjects were given recognition tests for sentences in the stories they had just read. Three classes of sentences were tested: OLD sentences, which had previously appeared in a story, NEW sentences, which were made up of invalid combinations of facts from the stories, and the INFERENCES described above. Subjects were instructed to use a response line below each sentence to answer two questions: (a) Did this exact sentence appear in a story you have read (OLD or NEW)? and (b) If not, is the sentence logically true with respect to the facts presented in the stories (TRUE or FALSE)? Thus, correct responses for the three types of items were OLD for old items, NEW and FALSE for new items, and NEW and TRUE for inference items.

There were twelve test sentences -- four of each of the three types of items -- for each set of stories. The test sentences were

blocked by story set, and these blocks were presented in the same order as the story sets to control any short-term retention of the stories. Subjects were not allowed to look back at the stories at any time during the testing.

Inference judgments. After a twenty minute delay (filled with a task unrelated to this experiment), subjects were given a randomized list of the sixteen inferences and sixteen new sentences from the preceding test. The subjects were instructed to circle TRUE to statements which were true (although not necessarily explicitly present in the previous stories) and FALSE to sentences which were either false or indeterminable. These responses were made on a response line below each sentence, as was a rating of the subject's confidence in his response on a 1 to 5 scale (5 = high confidence). Again, subjects were not allowed to look back at the stories at any time.

Syllogism judgments. Following the inference judgment task, subjects were presented 16 syllogisms which would prove the true inferences, and 16 false syllogisms involving the false inferences tested above. Subjects were given instructions on how to solve syllogisms, and were told to make true/false and confidence ratings to the syllogisms in the same way that they responded to the inference judgments above. This task concluded the experiment.

RESULTS AND DISCUSSION

Inference generation. The inferences generated by each subject in the Control and Fact conditions were examined for the sixteen

inferences around which the stories were designed. (These will be referred to as the critical inferences of the experiment.) Fact subjects generated more of these critical inferences than did Control subjects (probability of generation: Control = .40, Fact = .84; $t(22) = 5.14$, $p < .01$), although a significant number of inferences were missed by both groups ($.84 < 1.0$, $t(11) = 3.73$, $p < .001$). Thus, while it was possible for the Control subjects to generate some of the inferences on their own, the hints given to Fact subjects greatly increased the number of inferences they could locate and generate.

Immediate recognition test. The variable of interest is the probability of a subject's calling an OLD, NEW, or INFERENCE sentence old or new and true. While many interesting questions exist concerning subjects' abilities to discriminate sentences they have actually seen before from new and valid statements (Bransford and Franks, 1971), our concern here is whether or not subjects think a statement is true, regardless of their reasons. Failure to mark a true inference as "true" could be due to either of two possibilities. First, subjects may have failed to integrate the related sentences underlying an inference. Alternatively, they may have integrated the facts, but rejected the inference as not logically following from them. Since this chapter is primarily concerned with the subjects' memory capabilities, performance on the syllogism task was used to control performance on the inference task for such reasoning failures. Any inference whose corresponding syllogism was incorrect was ignored. Thus, the analyses below include only those inferences which the subject believed to be logically valid when all the necessary information was readily available. Performance on neither the true

nor the false syllogisms differed between groups (p(correct true syllogism: Control = .92, Fact = .92, Fact-and-Inference = .96; F(2,33) = 1.27, MSe = .006; p(correct false syllogism: Control = .80, Fact = .77, Fact-and-Inference = .68; F(2,33) = .82, MSe = .060), so this conditionalizing did not lead to any item selection effects.

The probability of a correct response was tested for a group (Control / Fact / Fact-and-Inference) by item (OLD / NEW / INFERENCE) analysis of variance. The interaction between group and item was significant, F(4,66) = 5.05, p < .001), as illustrated in Table 1. All three experimental groups called OLD items "true" equally often (F(2,33) < 1.0). However, Newman-Keuls tests (Winer, 1962) found that Control subjects called INFERENCES "true" significantly less often than Fact or Fact-and-Inference subjects (p < .01), and Control and Fact-and-Inference subjects correctly called NEW items "new" less often than did Fact subjects (p < .05). Further, only the Fact subjects performed better than chance on these NEW items (99 percent confidence intervals: Control = (.58, .41), SE = .038; Fact = (.70, .57), SE = .031; Fact-and-Inference = (.66, .43), SE = .052). In addition to the above interaction, the main effects of group (F(2,33) = 8.87, p < .001) and item (F(2,66) = 240.69, p < .001) were significant.

These data indicate that the Fact annotation resulted in the best performance of all three groups, as measured by the superiority of Fact subjects on both true inferences and false distractors. The poor performance of Fact-and-Inference subjects on NEW items confirms the earlier prediction that these subjects might read the stories without applying sufficient processing to insure their complete comprehension

and representation in memory. Therefore, we conclude that Fact annotation produces better knowledge of inferences than does Fact-and-Inference annotation, even on an immediate test.

Table 1

Probability of a correct response to OLD, NEW, and INFERENCE items on the initial recognition test. p(TRUE) = p(OLD) + p(NEW and TRUE).

	Control	Fact	Fact and Inference
p(TRUE¦OLD)	.96	.96	.96
p(TRUE¦INFERENCE)	.80	.96	.97
p(NEW and FALSE¦NEW)	.49	.63	.54

The importance of actually generating inferences during study can be seen by conditionalizing the probability of calling an inference "true" on whether or not the subject generated that inference (Table 2) during study. Such an analysis on the data from the Control and Fact subjects clearly shows that the most accurate predictor of performance on this task is the generation of inferences ($F(1,17) = 19.99$, $p < .001$). Neither the group effect nor the group by production interaction were significant (both F's < 1.0). The difference between the overall performance by Control and Fact subjects must then be due to the failure of Control subjects to produce as many inferences during study as did Fact subjects. The Control and Fact subjects' data for produced inferences did not differ from the overall Fact-and-Inference subjects' data ($F(2,28) < 1.0$), indicating that subjects' memory for a generated inference was at

least as good as for an explicitly stated fact.

Table 2

Probability of calling an inference TRUE, conditionalized on the generation of that inference: Immediate test.

	Control	Fact	Fact and Inference
Produced	.96	.98	.97
Not Produced	.66	.75	---

Another difficulty exists with Fact-and-Inference annotation. Providing overly specific "hints" about possible inferences may result in subjects' restriction of their searches to only those inferences for which suggestions exist. The inference generation data for this experiment were analyzed for valid inferences other than the critical inferences. Noncritical inference generation decreased as the amount of generation information increased (number of noncritical inferences generated: Control = 15.75, Fact = 13.00, Fact-and-Inference = 10.42). While this effect was not significant, due to the large between subject variance ($F(2,33) = 1.09$, $MSe = 78.10$), it deserves consideration. If subjects do, in fact, generate fewer inferences in the Fact-and-Inference condition, that finding would argue strongly in favor of Fact, rather than Fact-and-Inference annotation. The significance of noncritical inferences is especially great since, in many text situations, a clear set of desired inferences is not always available. Fact annotation would maximize the retention of those inferences that could be specified, while not severely restricting the

reader's independent generation of unspecified inferences. The advantage of Control subjects in generating noncritical inferences is outweighed by their poor performance on the critical inferences. Thus, Fact annotation appears to produce the best overall performance on critical and non-critical inferences.

Final inference and syllogism judgments. As in the initial recognition test, performance on the final inference task was conditionalized upon a correct response to the syllogism corresponding to each inference. Performance on this inference task is shown in Table 3. The groups did not differ statistically in rejecting false inferences ($F(2,33) = .51$, $MSe = .049$), but Control subjects were significantly poorer at verifying true inferences ($F(2,33) = 19.91$, $p < .001$, $MSe = .004$) than the other groups. Notice that the slight, although nonsignificant decrease in Fact-and-Inference subjects' accuracy at rejecting false inferences casts doubt upon the equivalence of Fact and Fact-and-Inference subjects' acceptance of true inferences. A true advantage to Fact annotation is instead suggested.

Table 3

Probability of correct inference judgment: Delayed test

	Control	Fact	Fact and Inference
p(true|true)	.82	.97	.97
p(false|false)	.63	.72	.68

If inference performance is conditionalized upon prior generation of the inference, the data shown in Table 4 are obtained. Again, the critical determinant of performance is clearly whether or not the inference was produced during study: only the generation effect of the analysis of variance ($F(1,17) = 8.12$, $p < .025$, $MSe = .033$) was significant. A similar finding exists in the confidence rating data shown in Table 4. Subjects are much more confident that an inference is true when it had been generated earlier: only the generation factor was significant ($F(1,17) = 7.80$, $p < .025$, $MSe = .445$). As before, the Control and Fact subjects' accuracies and confidence ratings for produced inferences did not differ from the overall Fact-and-Inference subjects' accuracies and confidence ratings (both analyses: $F(2,28) < 1.0$).

Table 4

Probability of verifying a true inference and confidence ratings, conditionalized on the generation of that inference: Delayed test.

p(correct)

	Control	Fact	Fact and Inference
Produced	.95	.97	.97
Not Produced	.73	.85	---

Confidence ratings
(+5 = certain true; -5 = certain false)

	Control	Fact	Fact and Inference
Produced	3.97	4.00	4.28
Not Produced	2.05	2.81	---

CONCLUSIONS

The primary conclusion to be drawn from these data is that inference generation is essential to long-term retention. This is evident from the analysis of inference verification conditionalized upon prior generation of inferences. Knowledge of generated inferences was near perfect, while knowledge of ungenerated inferences was only slightly above chance; this difference was maintained in both immediate and delayed testing. We can therefore conclude that, to insure maximum text retention, readers must be urged to view text reading as a process that requires integrative processing of the information they are reading while they are reading it.

To some extent, simply giving general instructions which show how to generate inferences and which stress the importance of such integration can produce significant acquisition: our control subjects were able to generate 40% of the possible critical inferences without any hints at all. This was far above the performance of subjects in the unpublished research discussed earlier. The goals and strategies of the reader during acquisition are thus a critical factor in the reading process. This goal orientation is the responsibility of the readers or the instructor in a classroom situation.

The responsibility of the text writer is constructing a text in a way that facilitates the inference generation process. The present results demonstrate the value of annotating texts to indicate related information that might be integrated. Of the two annotation methods considered here, we feel that the Fact annotation is better, for both theoretical and practical reasons. The extra information given to the

Fact and Inference subjects actually hindered them: Fact and Inference subjects were less able than Fact subjects to reject false statements in nearly all the measures studied here, sometimes performing no better than chance. These results, combined with the depressed generation of noncritical inferences, suggests that the Fact and Inference subjects relied too heavily upon this annotation for their understanding of the text, and did not engage in the semantic processing of the text that led to the superior performance of subjects with Fact annotation.

This latter point is especially important since what should ideally be taught in most educational situations is not simply a mass of facts, but also techniques by which these facts can be organized into a unified structure. As mentioned before, it is usually impractical (if not impossible) to specify every inference that could be derived from a text. The Fact manipulation thus has the advantage that it can cue a very specific inference that the text author might want made, while retaining the ability to indicate very general inferences. If some previously read material is generally relevant to the integration of a later portion of text, the Fact manipulation can reinstate this information without limiting the number and type of inferences that will be generated. Modeling the text after the Fact and Inference manipulation would likely prevent the generation of potentially important unspecified inferences.

REFERENCES

Bransford, J. and Franks, J. The abstraction of linguistic ideas. Cognitive Psychology, 1971, 2, 331-350.

Craik, F. I. M., and Lockhart, R. S. Levels of processing: A framework for memory research. Journal of Verbal Learning and Verbal Behavior, 1972, 11, 671-684.

Frase, L. T. Structural analysis of the knowledge that results from thinking about text. Journal of Educational Psychology Monograph, 1969, 60(6,Pt.2).

Frase, L. T. Integration of written text Journal of Educational Psychology, 1973, 65, 252-261.

Frase, L. T. Prose processing. In G. H. Bower (Ed.), The Psychology of Learning and Motivation. Vol. 9. New York: Academic Press, 1975.

Green, D. M., and Swets, J. A. Signal Detection Theory and Psychophysics. Huntington, New York: Krieger Publishing Co., 1966.

Haviland, S. and Clark, H. What's new? Acquiring new information as a process in comprehension. Journal of Verbal Learning and Verbal Behavior, 1974, 13, 512-521.

Hayes-Roth, B. Evolution of cognitive structures and processes. Psychological Review, 1977, 84, 260-278.

Hayes-Roth, B., and Thorndyke, P. W. Human processing of knowledge from text: Acquisition, integration, and reasoning (Chapter III). Report R-2256-ARPA, forthcoming, Rand Corporation, Santa Monica, California.

Hyde, T. S., and Jenkins, J. J. Differential effects of incidental tasks on the organization of recall of a list of highly associated words. Journal of Experimental Psychology, 1969, 82, 472-481.

Winer, B. J. Statistical principles in experimental design. New York: McGraw-Hill, 1962.